

REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-02-

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Project, Washington, DC 20503.

0450

ring
tion

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE	3. REPORT TYPE AND DATES COVERED 01 APR 01 TO 31 MAR 02
4. TITLE AND SUBTITLE OPCODE (ORLANDO PARALLEL COMPUTATION DEVELOPMENT ENVIRONMENT)			5. FUNDING NUMBERS F49620-01-1-0195
6. AUTHOR(S) Prof. Schiavone			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Central Florida Institute for Simulation and Training 3280 Progress Drive Orlando, FL 32826			8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 4015 Wilson Blvd, Room 713 Arlington, VA 22203-1954			10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-01-1-0195
11. SUPPLEMENTARY NOTES			
12a. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE, DISTRIBUTION UNLIMITED			12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 words) The OPCODE project has resulted in purchase of components for the construction of two 192-processor computing clusters, one located at 1ST (OPCODE I) and the other at SEECS UCF (OPCODE II). OPCODE I, shown in Figure 1, is completed and has been benchmarked using SCALAPACK at 86.5 GFLOPS using 144 nodes of the total 192 compute nodes (See Appendix B for results). The OPCODE II cluster is currently under construction, and is about half complete. These clusters employ a fast Ethernet network, linked together over a high-speed, fully stacked switch. Each node consists dual AMD 1500+ processors on a Tyan "Tiger" K7 motherboard, with 1 GB of DDR RAM, and 20 GB local disk storage per node. The compute nodes are rack mounted with two nodes per tray (see Figure 2). An alternative IEEE 1394 network is also be available on 44 of the nodes, on each machine. We have contributed extensively to developing IEEE 1394 drivers for Linux, and one of our patches is included in the Linux 2.4.x kernel.			
14. SUBJECT TERMS		20030115 088	
17. SECURITY CLASSIFICATION OF REPORT		18. SECURITY CLASSIFICATION OF THIS PAGE	15. NUMBER OF PAGES 30 16. PRICE CODE
19. SECURITY CLASSIFICATION OF ABSTRACT		20. LIMITATION OF ABSTRACT	

Heckler
NM

Final Report
**OPCODE (Orlando Parallel Computation
Development Environment)**

**Defense University Research Instrumentation
Program (DURIP)**

UCF Account Number: 64-02-517/16-23-915

Agency Name: AFOSR

Agency Number: F496200110195



**INSTITUTE for
SIMULATION
& TRAINING**

**University of Central Florida
Institute for Simulation and Training
3280 Progress Drive
Orlando, FL 32826
(407) 658-5000**

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

OPCODE

(Orlando Parallel Computation Development Environment)

Final Report

1. Introduction

This report describes research and support work performed by the University of Central Florida (UCF) and Institute for Simulation and Training (IST) in support of the OPCODE Project funded by the Defense University Research Instrumentation Program (DURIP).

The objective of the OPCODE project was to construct a large Beowulf-type computing cluster to support AFOSR and other DoD projects, as well as to support educational objectives in the areas of parallel and distributed computing. The original proposed cluster was to be composed of 192 Athlon 800 MHz processors, supported by 256 MB of RAM and 10 GB of disk storage per node. Due to AMD donations and price drops, as well as the choice to use a dual-processor motherboard design, we were able to purchase sufficient parts to build two 192 processor clusters, using AMD 1500+ processors running at 1.4 Ghz.

1.1 Background

Beowulf is a multi-computer architecture used for parallel computations. It is a system that usually consists of one server node, and one or more client nodes connected together via Ethernet or some other network. It is a system built using commodity hardware components, such as any PC that is capable of running Linux, standard Ethernet adapters, and switches. It does not contain any custom hardware components and it is trivially reproducible. Beowulf also uses commodity software like Parallel Virtual Machine (PVM) or Message Passing Interface (MPI). The server node controls the whole cluster and serves files to the client nodes. It is also the cluster's console and the gateway to the outside world. Client nodes in a Beowulf system are dumb: They are configured and controlled by the server node, and do only what they are told to

Beowulf systems have been constructed from a variety of parts. For the sake of performance some non-commodity components (i.e. produced by a single manufacturer) occasionally have been employed in a few implementations.

A CLASS I Beowulf cluster is built entirely from commodity "off-the-shelf" parts. The advantages of a CLASS I system are:

- hardware is available from multiple sources (low prices, easy maintenance)
- no reliance on a single hardware vendor
- free driver support from Linux community
- hardware components based on standards (SCSI, Ethernet, etc.)

In the taxonomy of parallel computers, Beowulf clusters fall somewhere between MPP (Massively Parallel Processors, like the Convex SPP, Cray T3D, Cray T3E, CM5, etc.) and NOWs (Networks of Workstations). A Beowulf cluster benefits from developments in both these classes of architecture. MPP's are typically larger and have a lower latency interconnect network than Beowulf clusters. Using MPPs, programmers are required to consider locality, load balancing, granularity, and communication overheads. Most programs for MPPs are developed in message passing style. Such programs can be readily ported to Beowulf clusters.

Programming a NOW is usually an attempt to harvest unused cycles on an already installed base of workstations in a lab or on a campus. Programming in this environment requires algorithms that are extremely tolerant of load balancing problems and large communication latency. These programs will directly run on a Beowulf. A Beowulf class cluster computer is distinguished from a NOW by several subtle but significant characteristics. First, the nodes in the cluster are dedicated to the cluster. This helps ease the load balancing problem, because the performance of individual nodes is not subject to external factors. Also, since the interconnection network is isolated from the external network, the network load is determined only by the application being run on the cluster. This eases the problems associated with unpredictable latency in NOWs. All the nodes in the cluster are within the administrative jurisdiction of the cluster. For example, the Beowulf software provides a global process ID that enables a mechanism for a process to send signals to a process on another node of the system. This is not allowed on a NOWs. Finally, operating system parameters can be tuned to improve performance. For example, a workstation should be tuned to provide the interactive feel (instantaneous responses, short buffers, etc), but in cluster the node can be tuned to provide better throughput for coarser grain jobs because they are not interacting with users.

In 1999, IST built a CLASS I Beowulf cluster, consisting of 17 nodes and 34 Pentium II processors, with 256 Mb of RAM and 8.6 Gb of local storage per node. Cluster communications take place over a three channel-bonded fast ethernet network, using low-latency switches to minimize network traffic. An IEEE 1394 standard alternate network has also been installed to offer lower latency, improved bandwidth, and deterministic communications.

In 2000, a 128 node cluster was constructed, located in the School of Electrical Engineering and Computer Science (SEECs) at UCF. This cluster, "SCEROLA", consists of 128 900 Mhz Athlon processors, with 32 Gb total RAM and 1.6 Terabytes total secondary storage.

2 Technical Approach

Based on our experience in building previous clusters, our approach was to custom-design a cluster based on the most currently available commodity hardware that yielded the best cost/performance ratio. We investigated the cost/performance ratio of "prebuilt" Beowulf clusters that are available commercially, and determined that we could build a cluster with between 2-3 times increase in the performance by designing and building our own cluster from individual components. As an additional benefit, we also viewed the exercise of cluster design

and construction as a worthy academic and educational experience. The tradeoff made for these benefits was that financial support for the considerable labor involved in the design, construction and installation of the cluster was not funded by the DURIP program, and had to be supported by other means. Fortunately, some labor support was provided by the Advanced Tactical Engagement Simulation Science and Technology (A-TES STO) program funded by U.S. Army STRICOM, as was mentioned in the original proposal to AFOSR. Other financial support for labor was derived from UCF matching funds. Ultimately, however, the lack of sufficient labor funding has resulted in delays in construction that were not fully anticipated in the original schedule.

2.1 Current Status

The OPCODE project has resulted in purchase of components for the construction of two 192-processor computing clusters, one located at IST (OPCODE I) and the other at SEECs UCF (OPCODE II). OPCODE I, shown in Figure 1, is completed and has been benchmarked using SCALAPACK at 86.5 GFLOPS using 144 nodes of the total 192 compute nodes (See Appendix B for results). The OPCODE II cluster is currently under construction, and is about half complete. These clusters employ a fast Ethernet network, linked together over a high-speed, fully stacked switch. Each node consists dual AMD 1500+ processors on a Tyan "Tiger" K7 motherboard, with 1 GB of DDR RAM, and 20 GB local disk storage per node. The compute nodes are rack mounted with two nodes per tray (see Figure 2). An alternative IEEE 1394 network is also be available on 44of the nodes, on each machine. We have contributed extensively to developing IEEE 1394 drivers for Linux, and one of our patches is included in the Linux 2.4.x kernel.




Figure 1 OPCODE I cluster at IST




Figure 2 – Two compute nodes from OPCODE I showing tray for rack-mounting.

2.2 Chronology of events and design decisions

01-2001 – Notification of award

02-2001 – Meetings held with UCF Office of Research to determine level of UCF Matching funds. Matching funds totaling \$54, 250 were awarded under the Florida State High-Tech Corridor I-4 Program. Additional funding from the UCF Presidential Equipment Program was not made available.

03-2001 – It was decided that two clusters would be built, and the decision was made to split the award with half administered by IST and the other half by the UCF School of Electrical and Computer Engineering (SEECs).

04-2001 – Funds awarded and accounts provided for both AFOSR and UCF matching funds. Due to UCF funding procedures, nearly all of the matching funds were deferred until after the end of the UCF fiscal year (ending 06-2001)

05-2001 - A location for the IST cluster, OPCODE I, was determined, and cooling requirements were estimated. Quotes for an air conditioning upgrade with a 6.5 ton capacity were obtained. The online Beowulf site at http://www.phy.duke.edu/brahma/beowulf_online_book/node43.html was helpful in determining the cooling requirements.

06-2001 - Additional quotes for the power and air conditioning were obtained, allowing for 300 Amps of power. Quotes were obtained for basic items that were not particular to the final design decisions, such as for RAM and NICs.

07-2001 - Performance comparisons between the currently available AMD processors and the Intel P4 were considered. A decision was made to use the AMD processors in the cluster design, based on the following considerations: 1) The ability of the AMD processors to perform 3 floating point operations per clock cycle (vs. 2 floating point operations for the P4), 2) The superior cost/performance ratio offered by the AMD processors, and 3) The high cost of Rambus memory modules required by the P4. Two useful sources of information that were considered follow:

<http://arstechnica.com/cpu/01q2/p4andg4e/p4andg4e-1.html>

<http://www.emulators.com/pentium4.htm>

08-2001 Elitegroup K755A motherboards with SiS 736 chipsets were quoted for \$ 65.00. These boards come with onboard Ethernet, but since the Ethernet capability on these boards was based on the RealTek chip (RTL8201L-Chip) they were ruled out, as they are slower than NICs based on the Tulip or Intel chipsets because they do not do dma on receive, resulting in an extra mem copy.

The HP procurve 4100gl switch was investigated due to its low cost, although it does not allow for channel bonding.

09-2001 - A system based on the dual processor Tyan Tiger motherboards using Athlon 1.2 Ghz Palomino processors with 1 GB of PC2100 DDR memory was purchased for testing purposes. A design based on these dual-processor motherboards was chosen. Power requirements were tested by measuring current draw under a variety of conditions, as shown below:

<u>Situation</u>	<u>Current (Amps)</u>
Not on	0.05
powerup	1.20
post	1.15
fsck in linux	1.11
Linux on but idle	1.08
distributed.net rc5-64	1.33
rc5-64 & updatedb	1.36
cpuburn	1.18
rc5-64 & cpuburn	1.29

It was noted that these figures were lower than expected with regards to cpuburn vs. rc5-64, the reason being that rc5-64 is multi-threaded and so was using both processors, while cpuburn was single threaded so did not use the second processor.

System stability was observed to be good.

09-2001 Carrie Hermann, Manager of AMD University Relations, was contacted regarding possible processor donations. A proposal was submitted to AMD requesting 192 Athlon 4 "Palamino" processors. This proposal is attached as Appendix C

10-2001 – AMD proposal denied, but referral is made to AMD University Parts Program. Racks purchased for OPCODE I cluster.

11-2001 - Performance of the raid arrays that will be used for the servers in the new cluster is tested. Performance was tested for both software and hardware raid levels 0 and 5. We also tested the performance with different file systems. Best performance is noted using software raid and XFS file systems.

The power and air conditioning for OPCODE I at IST are completed. With a minimal heat load in the room the temperature stays at a constant 60F. Purchased 5 racks and some networking supplies to be used in OPCODE II.

12-2001 – A decision was made to purchase the Extreme Networks Alpine 3804 switch instead of the HP Procurve Switch that had previously been considered. This decision was based on information that ProCurve 4108GL will not do wire speed to all ports. It is rated at 36.6 Gbps fabric speed, compared with 38.4 Gbps with all of the ports used at 100TX. Now granted this number does not apply to us because we will not be using all of the ports. Customer reports indicated that the previous HP Procurve model would only perform at about 2/3 of its rated speed. HP specifies that the switch has a latency of <10us (FIFO) and is able to have up to 6 trunks with 4 ports in each trunk. Reports on the Beowulf.org list indicated that some people had a hard time getting trunking (bonding) to work on the previous model Procurve switch.

The Extreme Networks Alpine 3804 chassis is rated at 32 Gbps which is more than can be plugged into 128 10/100 ports the chassis can handle (Alpine 3808 can handle 64 Gbps with 256 ports). There have been independent tests that confirm that the switch will do wire speed to all ports with no blocking. Customer reported <5us latency for the switch, although the specification is not given by Extreme Networks. The Alpine switch can have as many trunks on the switch as needed, with at least 4 ports per trunk. An Extreme Networking representative informed us that they have an upgrade process, allowing for an possible future upgrade to channel bonding to all of the nodes by replacing the chassis and purchasing a couple of ethernet modules. The smaller chassis we chose does not allow for channel bonding to every node.

01-2002 – Purchased the Alpine switches. Received donation from AMD of 60 AMD 1500+ processors, and took bids on the remaining 176 AMD 1500+ processors required for each cluster.

02-2002 The motherboards and processors were ordered. The racks being used for the ISTOPCODE cluster were painted and supplies were purchased for making the rear panels. One of the Alpine Extreme switches was tested by using it temporarily in the SECS SCEROLA cluster.

03-2002 We finished purchasing most of the components for the OPCODE clusters. There are still a couple of items that need to be purchased on some of the matching funds accounts. Wiring the network of the cluster was begun. Custom cables were made for use between the patch panel and the switch, and cables were run from the patch panel to each of the computer trays. The trays that will be holding the motherboards were designed.

04-2002 Finished wiring the switch to the patch panel and started to configure the switch for use on the OPCODE I cluster and IST network. Two vlans were setup on the switch, one for the cluster network and one for the IST network. Some computers are to be on both vlans so modifications of the network stacks on those machines were necessary to allow for the larger packets that tagged vlans impose. The design for the trays that the motherboards will mount on was finalized and has been sent of to the manufacture so that a prototype can be made. The frontend machine for the OPCODE I cluster has been setup and is being configured.

05-2002 – Present

The OPCODE I cluster is completed and running. SCALPACK benchmarks give a top performance rating of approximately 86.5 GFLOPS running 144 of the nodes, as shown in Appendix B. OPCODE II is under construction, with about half of the trays finished at present.

Appendix A – Financial Report

Table I – Original Submitted Budget

(Note – Additional UCF Match ultimately was not obtained, by decision of UCF Office of Research)

"OPCODE (ORLANDO PARALLEL COMPUTATION DEVELOPMENT ENVIRONMENT)"

DIRECT COSTS	SPONSOR COST	UCF MATCH (1)	TOTAL BUDGET	ADDITIONAL UCF MATCH (2)
LABOR	\$ -	\$ -	\$ -	\$ -
EQUIPMENT	\$ 217,000	\$ 54,250.00	\$ 271,250	\$ 217,000
TRAVEL	\$ -	\$ -	\$ -	\$ -
MISCELLANEOUS MATERIALS AND SUPPLIES	\$ -	\$ -	\$ -	\$ -
TOTAL DIRECT COSTS	\$ 217,000	\$ 54,250	\$ 271,250	\$ 217,000
INDIRECT COST (Total Direct Costs Less Equipment x 42.5%)	\$ -	\$ -	\$ -	\$ -
TOTAL COST	<u>\$ 217,000</u>	<u>\$ 54,250</u>	<u>\$ 271,250</u>	<u>\$ 217,000</u>

NOTES:

(1) "UCF MATCH" column represents matching funds to be committed to this funding.

(2) "ADDITIONAL UCF MATCH" column represents matching funds that may be offered pending University approval.

Table II - Original Proposed Equipment List (prices obtained 08/16/2000)

Item	Description	Source	Cost	Qty	Total Cost
ASUS K7V VIA KX133	5 PCI, No Audio, - 1Ghz	www.necxdirect.com	131.95	192	25334.40
48"x18"x72" black wire- frame shelves		Local Hardware Store (Lowe's or equivalent)	64.00	12	768
AMD Athlon Thunderbird	800Mhz Socket A	www.motherboards.com	196.00	192	37632
Hawking 15' Cat.5e cable 5- packs	15' transparent color-coded	Buy.Com	12.00	80	960
Quantum 10.2GB EIDE	Ultra DMA66/33,7200 RPM	McGlen, www.mcglenn.com	101.72	192	19530.24
SIIG IEEE1394 3- port		www.necxdirect.com	49.95	192	9590.4
30GB FIREWIRE EXTERNAL HD	5200RPM, 16.6MB per second data transfer	www.necxdirect.com			369.95
PC133 SDRAM 8NS 256MB		www.mwave.com	305.00	384	117120
Misc					3000
SUPERSTAC K II SWITCH 3300 MM	www.3com.com/pro ducts/switches/super stack/ss2 3300mm	McGlen, www.mcglenn.com	1,625.62	24	39014.88
Smartlink 100Mb/s NIC 10-packs	www.mwave.com	www.mwave.com	80.00	60	4800
MID ATX 6- BAY 230W +FLOPPY		www.necxdirect.com	67.95	192	13046.4
TOTAL					271166.27

Table III - OPCODE I Machine Components and Budget

Nodes	96			
Spares	2			
Total Nodes	98			
Servers	4			
Frontend	1			
Total	103			
		Cluster Node		
Item	Vendor	Description	Count	Cost
Processor	D&H	Athlon MP 1500+ 1 year Warranty	2	\$ 159.9
Fan	Monarch	CoolerMaster DP5-I11A-A1	2	\$ 7.0
Motherboard	MicroPro	TYAN s2460 AMD 760MP 266/200 FSB 4 DIMMS DDR 1/5	1	\$ 175.0
Memory	Muskin	Muskin PC2100 256 MB CL 2.5 ECC Registered	2	\$ 42.0
NIC	AMASTORE	Netgear FA310TX	1	\$ 10.5
Hard drive	TC	Maxtor 30GB ATA 100 5400RPM 2 MB	1	\$ 85.0
power supply	Axion Tech	Dynapower DP350A 350 Watt 12V atx power supply	1	\$ 42.2
			Total	
			Node Total	
		Server Node		
Item	Vendor	Description	Count	Cost
Processor	D&H	Athlon MP 1500+ 1 year Warranty	2	\$ 159.9
Fan	Monarch	CoolerMaster DP5-I11A-A1	2	\$ 7.0
Motherboard	Sybercom	TYAN s2460 AMD 760MP 266/200 FSB 4 DIMMS DDR 1/5	1	\$ 175.0
Memory	Muskin	Muskin PC2100 256 MB CL 2.5 ECC Registered	4	\$ 42.0

NIC	AMASTORE	D-Link DFE-570TX 4 port Ethernet card	1	\$ 109.0
power supply	Axion Tech	Dynapower DP400 400 Watt 12V atx power supply	1	\$ 62.0
server monitor				
server harddrives	Hyper Micro	Maxtor 81GB ATA 100 5400RPM 2 MB	4	\$ 155.0
server raid	Hyper Micro	3ware 4-port ata/66 RAID 5 controller (6410)	1	\$ 179.0
server case			1	
			Total	
			Server Total	
			Computer Total	
		Extreme Switch		
Item	Vendor	Description	Count	Cost
switch	Extreme Networks	Alpine 3804 Chassis	1	\$ 4,430.0
switch	Extreme Networks	Alpine SMMi Basic L3	1	\$ 3,385.5
switch	Extreme Networks	Alpine 3800 FM-32Ti (32port 10/100Tx module)	4	\$ 1,116.5
switch	Extreme Networks	Alpine 3800 PS	1	\$ 1,396.5
switch	Extreme Networks	Alpine 3804 Service contract (48 Hours)	1	\$ 1,322.0
			Total	
		Networking		
Item	Vendor	Description	Count	Cost
networking	Sofisticated	cat5 2xRJ45 surface mount box	114	\$ 4.0
networking	Skycraft	48 Port cat 5 patch panel	5	\$ 48.0
networking	Home Depot	Telemaster RJ11/RJ45 crimper Ratcheting	1	\$ 42.5
networking	Sofisticated	Cat 5 RJ45 jacks solid and stranded 1000piece	1	\$ 59.0
networking	PCTek Online	Cat5e boots	500	\$ 0.1
networking	PCTek Online	1000' Cat5e Solid Cable 350Mhz Yellow	4	\$ 44.0
			Total	

			Networking Total	
		Rack Components		
Item	Vendor	Description	Count	Cost
Chassis	Turf Equipment	Motherboard Trays w/ Power supply brackets	53	\$ 29.4
Chassis	Turf Equipment	Motherboard Tray rails	53	\$ 3.0
Chassis	Turf Equipment	Rail Supports	38	\$ 10.0
Chassis	Turf Equipment	Stop brackets	14	\$ 3.0
			Total	
		Misc		
Item	Vendor	Description	Count	Cost
Air Flow	Home Depot	Door bottom molding	1	\$ 6.9
Tools	Home Depot	Versapack Cordless Screwdriver	1	\$ 28.9
fasteners	Skycraft	500 4" cable ties and 1000 6" cable ties	1	\$ 27.4
fasteners	Skycraft	Standoffs	500	\$ 0.0
fasteners	Skycraft	computer screws #6-32 1/4	3000	\$ 0.0
Fans	Home Depot	Lasko 20" box fan	7	\$ 10.9
Fans	Skycraft	19" chassis fan	1	\$ 100.0
Air Flow	Home Depot	Hardboard	5	\$ 5.4
Paint	Home Depot	Rustoleum Hard Hat Spray Paint (Black)	10	\$ 4.9
Power	Home Depot	Belkin 7 outlet 15A 12ft cord 752J \$20000 surge protector	18	\$ 18.9
UPS	Warehouse	Tripp Lite Smart Pro Net 2200VA	1	\$ 715.0
UPS	Warehouse	Tripp Lite Smart Pro 1400XL	2	\$ 559.0
			Total	
		KVM		

Item	Vendor	Description	Count	Cost
KVM	Micro Warehouse	Belkin OMNIVIEW Matrix2 2x8 port KVM	1	\$ 618.00
KVM	Micro Warehouse	OMNIVIEW CAT5 Extender	2	\$ 222.00
KVM	Micro Warehouse	25ft Matrix Cable PS2	4	\$ 41.00
KVM	Micro Warehouse	6ft Matrix Cable PS2	4	\$ 20.00
			Total	
		Shipping		
Item	Vendor	Description	Count	Cost
shipping	PCTek Online	Cat5e and RJ45 Boots		
shipping	Sofisticated	cat5 2xRJ45 surface mount box and RJ45 plugs		
shipping	Muskin	Memory		
shipping	Hyper Micro	Server harddrives and Raid 5		
			Total	
		Donations		
Item	Vendor	Description	Count	Cost
Processors	AMD	Athlon MP 1500+ 1 year Warranty	30	\$ 159.00
			Total	
		Other Expenses		
Item	Vendor	Description	Count	Cost
Dev machines		Dual Processor Test and Front End Support Machines	15	\$ 900.00
			Grand Total	

Table IV OPCODE II Cluster Components (Identical to OPCODE I)

Nodes	96			
Spares	2			
Total Nodes	98			
Servers	4			
Frontend	1			
Total	103			
		Cluster Node		
Item	Vendor	Description	Count	Cost
Processor	D&H	Athlon MP 1500+ 1 year Warranty	2	\$ 159.9
Fan	Monarch	CoolerMaster DP5-I11A-A1	2	\$ 7.0
Motherboard	MicroPro	TYAN s2460 AMD 760MP 266/200 FSB 4 DIMMS DDR 1/5	1	\$ 175.0
Memory	Muskin	Muskin PC2100 256 MB CL 2.5 ECC Registered	2	\$ 42.0
NIC	AMASTORE	Netgear FA310TX	1	\$ 10.5
Hard drive	TC	Maxtor 30GB ATA 100 5400RPM 2 MB	1	\$ 85.0
power supply	Axion Tech	Dynapower DP350A 350 Watt 12V atx power supply	1	\$ 42.2
			Total	
			Node Total	
		Server Node		
Item	Vendor	Description	Count	Cost
Processor	D&H	Athlon MP 1500+ 1 year Warranty	2	\$ 159.9
Fan	Monarch	CoolerMaster DP5-I11A-A1	2	\$ 7.0
Motherboard	Sybercom	TYAN s2460 AMD 760MP 266/200 FSB 4 DIMMS DDR 1/5	1	\$ 175.0
Memory	Muskin	Muskin PC2100 256 MB CL 2.5 ECC Registered	4	\$ 42.0
NIC	AMASTORE	D-Link DFE-570TX 4 port Ethernet card	1	\$ 109.0
power supply	Axion Tech	Dynapower DP400 400 Watt 12V atx power supply	1	\$ 62.0
server monitor				
server harddrives	Hyper Micro	Maxtor 81GB ATA 100 5400RPM 2 MB	4	\$ 155.0

server raid	Hyper Micro	3ware 4-port ata/66 RAID 5 controller (6410)	1	\$ 179.0
server case			1	
			Total	
			Server Total	
			Computer Total	
		Extreme Switch		
Item	Vendor	Description	Count	Cost
switch	Extreme Networks	Alpine 3804 Chassis	1	\$ 4,430.0
switch	Extreme Networks	Alpine SMMi Basic L3	1	\$ 3,385.5
switch	Extreme Networks	Alpine 3800 FM-32Ti (32port 10/100Tx module)	4	\$ 1,116.5
switch	Extreme Networks	Alpine 3800 PS	1	\$ 1,396.5
switch	Extreme Networks	Alpine 3804 Service contract (48 Hours)	1	\$ 1,322.0
			Total	
		Networking		
Item	Vendor	Description	Count	Cost
networking	Sofisticated	cat5 2xRJ45 surface mount box	114	\$ 4.0
networking	Skycraft	48 Port cat 5 patch panel	5	\$ 48.0
networking	Home Depot	Telemaster RJ11/RJ45 crimper Ratcheting	1	\$ 42.5
networking	Sofisticated	Cat 5 RJ45 jacks solid and stranded 1000piece	1	\$ 59.0
networking	PCTek Online	Cat5e boots	500	\$ 0.1
networking	PCTek Online	1000' Cat5e Solid Cable 350Mhz Yellow	4	\$ 44.0
			Total	
			Networking Total	
		Rack Components		

Item	Vendor	Description	Count	Cost
Chassis	Turf Equipment	Motherboard Trays w/ Power supply brackets	53	\$ 29.5
Chassis	Turf Equipment	Motherboard Tray rails	53	\$ 3.0
Chassis	Turf Equipment	Rail Supports	38	\$ 10.0
Chassis	Turf Equipment	Stop brackets	14	\$ 3.0
			Total	
		Misc		
Item	Vendor	Description	Count	Cost
Air Flow	Home Depot	Door bottom molding	1	\$ 6.9
Tools	Home Depot	Versapack Cordless Screwdriver	1	\$ 28.9
fasteners	Skycraft	500 4" cable ties and 1000 6" cable ties	1	\$ 27.5
fasteners	Skycraft	Standoffs	500	\$ 0.0
fasteners	Skycraft	computer screws #6-32 1/4	3000	\$ 0.0
Fans	Home Depot	Lasko 20" box fan	7	\$ 10.9
Fans	Skycraft	19" chassis fan	1	\$ 100.0
Air Flow	Home Depot	Hardboard	5	\$ 5.5
Paint	Home Depot	Rustoleum Hard Hat Spray Paint (Black)	10	\$ 4.9
Power	Home Depot	Belkin 7 outlet 15A 12ft cord 752J \$20000 surge protector	18	\$ 18.9
UPS	Warehouse	Tripp Lite Smart Pro Net 2200VA	1	\$ 715.0
UPS	Warehouse	Tripp Lite Smart Pro 1400XL	2	\$ 559.0
			Total	
		KVM		
Item	Vendor	Description	Count	Cost
KVM	Micro Warehouse	Belkin OMNIVIEW Matrix2 2x8 port KVM	1	\$ 618.0
KVM	Micro Warehouse	OMNIVIEW CAT5 Extender	2	\$ 222.0

KVM	Micro Warehouse	25ft Matrix Cable PS2	4	\$ 41.0
KVM	Micro Warehouse	6ft Matrix Cable PS2	4	\$ 20.0
			Total	
		Shipping		
Item	Vendor	Description	Count	Cost
shipping	PCTek Online	Cat5e and RJ45 Boots		
shipping	Sofisticated	cat5 2xRJ45 surface mount box and RJ45 plugs		
shipping	Muskin	Memory		
shipping	Hyper Micro	Server harddrives and Raid 5		
			Total	
		Donations		
Item	Vendor	Description	Count	Cost
Processors	AMD	Athlon MP 1500+ 1 year Warranty	30	\$ 159.0
			Total	
		Other Expenses		
Item	Vendor	Description	Count	Cost
Dev machines		Dual Processor Test and Front End Support Machines	15	\$ 900.0
			Grand Total	

Table V IST Infrastructure Upgrades (from UCF Cost-sharing funds)

		Air and Power Infrastructure Upgrades (IST)			
Air & Power	Lincoln	300 Amp feed thru 84 circuit panel & 1 Service wire	1	\$10,070.00	\$ 10,070.00
Air & Power	Lincoln	6.5 ton HVAC unit & install	1	\$ 8,977.00	\$ 8,977.00
Air & Power	Lincoln	Construction Management Fee (5%)	1	\$ 952.35	\$ 952.35
Air & Power	Lincoln	30 amp 240 volt outlet	1	\$ 225.00	\$ 225.00
Air & Power	Lincoln	20 amp 120 volt outlet	14	\$ 125.00	\$ 1,750.00
			Total		\$ 21,974.35

Additional UCF Cost-sharing funds were expended on labor, and cluster computer laboratory hardware and software support,

Appendix B - Benchmark results for OPCODE I

```

HPLinpack benchmark input file
Innovative Computing Laboratory, University of Tennessee
HPL.out      output file name (if any)
6            device out (6=stdout,7=stderr,file)
1            # of problems sizes (N)
66000      Ns
2            # of NBs
60 64      NBs
1            # of process grids (P x Q)
12          Ps
12          Qs
16.0        threshold
3            # of panel fact
0 1 2       PFACTs (0=left, 1=Crout, 2=Right)
2            # of recursive stopping criterium
2 4         NBMINs (>= 1)
1            # of panels in recursion
2            NDIVs
3            # of recursive panel fact.
0 1 2       RFACTs (0=left, 1=Crout, 2=Right)
1            # of broadcast
3            BCASTs (0=1rg,1=1rM,2=2rg,3=2rM,4=Lng,5=LnM)
1            # of lookahead depth
0 1         DEPTHS (>=0)
2            SWAP (0=bin-exch,1=long,2=mix)
64          swapping threshold
0            L1 in (0=transposed,1=no-transposed) form
0            U in (0=transposed,1=no-transposed) form
1            Equilibration (0=no,1=yes)
8            memory alignment in double (> 0)

```

```

=====
HPLinpack 1.0 -- High-Performance Linpack benchmark -- September 27, 2000
Written by A. Petitet and R. Clint Whaley, Innovative Computing Labs., UTK
=====

```

An explanation of the input/output parameters follows:

```

T/V      : Wall time / encoded variant.
N        : The order of the coefficient matrix A.
NB       : The partitioning blocking factor.
P        : The number of process rows.
Q        : The number of process columns.
Time     : Time in seconds to solve the linear system.
Gflops   : Rate of execution for solving the linear system.

```

The following parameter values will be used:

```

N      : 66000
NB     : 60      64
P      : 12
Q      : 12
PFACT  : Left    Crout    Right
NBMIN  : 2       4

```



```

NDIV      :      2
RFACT     :      Left      Crout      Right
BCAST     :      2ringM
DEPTH     :      0
SWAP      :      Mix (threshold = 64)
L1        :      transposed form
U         :      transposed form
EQUIL     :      yes
ALIGN     :      8 double precision words

```

-
- The matrix A is randomly generated for each test.
 - The following scaled residual checks will be computed:
 - 1) $\|Ax-b\|_{\infty} / (\text{eps} * \|A\|_1 * N)$
 - 2) $\|Ax-b\|_{\infty} / (\text{eps} * \|A\|_1 * \|x\|_1)$
 - 3) $\|Ax-b\|_{\infty} / (\text{eps} * \|A\|_{\infty} * \|x\|_{\infty})$
 - The relative machine precision (eps) is taken to be 1.110223e-16
 - Computational tests pass if scaled residuals are less than 16.0

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
W03L2L2      66000    60     12     12          2297.85          8.341e+01

```

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) =      0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) =      0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) =      0.0015099 ..... PASSED
=====

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
W03L2L4      66000    60     12     12          2255.11          8.499e+01

```

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) =      0.0207939 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) =      0.0084675 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) =      0.0015126 ..... PASSED
=====

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
W03L2C2      66000    60     12     12          2263.99          8.466e+01

```

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) =      0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) =      0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) =      0.0015099 ..... PASSED
=====

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
W03L2C4      66000    60     12     12          2258.37          8.487e+01

```

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) =      0.0257147 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) =      0.0104713 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) =      0.0018706 ..... PASSED
=====

```

```

=====
T/V          N      NB      P      Q          Time          Gflops
-----
W03L2R2      66000    60     12     12          2249.42          8.521e+01

```

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03L2R4	66000	60	12	12	2259.19	8.484e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0211940 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0086304 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015418 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2L2	66000	60	12	12	2243.15	8.545e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2L4	66000	60	12	12	2253.61	8.505e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207939 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084675 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015126 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2C2	66000	60	12	12	2243.70	8.543e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2C4	66000	60	12	12	2254.58	8.501e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0257147 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0104713 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0018706 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2R2	66000	60	12	12	2250.39	8.517e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
W03C2R4	66000	60	12	12	2255.29	8.499e+01

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0211940 ..... PASSED

```

```

||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0086304 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015418 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2L2	66000	60	12	12	2255.07	8.500e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2L4	66000	60	12	12	2245.12	8.537e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207939 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084675 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015126 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2C2	66000	60	12	12	2248.07	8.526e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2C4	66000	60	12	12	2251.32	8.514e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0257147 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0104713 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0018706 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2R2	66000	60	12	12	2252.19	8.510e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0207560 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0084520 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015099 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03R2R4	66000	60	12	12	2252.21	8.510e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0211940 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0086304 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_oo * ||x||_oo ) = 0.0015418 ..... PASSED
=====

```

T/V	N	NB	P	Q	Time	Gflops
-----	---	----	---	---	------	--------

W03L2L2	66000	64	12	12	2231.95	8.588e+01
---------	-------	----	----	----	---------	-----------

```

||Ax-b||_oo / ( eps * ||A||_1 * N ) = 0.0217185 ..... PASSED
||Ax-b||_oo / ( eps * ||A||_1 * ||x||_1 ) = 0.0088440 ..... PASSED

```

||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015799 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03L2L4	66000	64	12	12	2238.98	8.561e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0217791 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0088687 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015843 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03L2C2	66000	64	12	12	2231.68	8.589e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0217185 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0088440 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015799 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03L2C4	66000	64	12	12	2237.69	8.566e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0233270 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0094990 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0016969 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03L2R2	66000	64	12	12	2240.78	8.554e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0217185 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0088440 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015799 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03L2R4	66000	64	12	12	2235.19	8.575e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0209615 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0085357 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015248 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03C2L2	66000	64	12	12	2252.16	8.511e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0217185 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0088440 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015799 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03C2L4	66000	64	12	12	2229.40	8.597e+01

||Ax-b||_oo / (eps * ||A||_1 * N) = 0.0217791 PASSED
 ||Ax-b||_oo / (eps * ||A||_1 * ||x||_1) = 0.0088687 PASSED
 ||Ax-b||_oo / (eps * ||A||_oo * ||x||_oo) = 0.0015843 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03C2C2	66000	64	12	12	2239.24	8.560e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217185 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088440 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015799 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03C2C4	66000	64	12	12	2239.84	8.557e+01
Ax-b _oo / (eps * A _1 * N) =						0.0233270 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0094990 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0016969 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03C2R2	66000	64	12	12	2237.44	8.567e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217185 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088440 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015799 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03C2R4	66000	64	12	12	2234.10	8.579e+01
Ax-b _oo / (eps * A _1 * N) =						0.0209615 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0085357 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015248 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03R2L2	66000	64	12	12	2233.78	8.581e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217185 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088440 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015799 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03R2L4	66000	64	12	12	2230.70	8.592e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217791 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088687 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015843 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03R2C2	66000	64	12	12	2240.57	8.555e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217185 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088440 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015799 PASSED

T/V	N	NB	P	Q	Time	Gflops
W03R2C4	66000	64	12	12	2234.19	8.579e+01
Ax-b _oo / (eps * A _1 * N) =						0.0233270 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0094990 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0016969 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03R2R2	66000	64	12	12	2227.54	8.605e+01
Ax-b _oo / (eps * A _1 * N) =						0.0217185 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0088440 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015799 PASSED
T/V	N	NB	P	Q	Time	Gflops
W03R2R4	66000	64	12	12	2239.26	8.560e+01
Ax-b _oo / (eps * A _1 * N) =						0.0209615 PASSED
Ax-b _oo / (eps * A _1 * x _1) =						0.0085357 PASSED
Ax-b _oo / (eps * A _oo * x _oo) =						0.0015248 PASSED

Finished 36 tests with the following results:
 36 tests completed and passed residual checks,
 0 tests completed and failed residual checks,
 0 tests skipped because of illegal input values.

End of Tests.

Explanation of Benchmark results:

All of the output lines of importance will look like this.

W03L2L2 66000 60 12 12 2297.85 8.341e+01

W: don't know
 0: depth
 3: type of broadcast
 L: RFACTs recursive panel factorization (Left, Crout, Right)
 2: NDIV
 L: PFACTs panel factorization (Left, Crout, Right)
 2: NBMIN recursion stops when the current panel is made of
 less than or equal to NBMIN columns
 66000: The order of the coefficient matrix A
 60: NB The partitioning blocking factor

12: P the number of process rows
12: Q the number of process columns
2297.85: Time in seconds
8.341e+01: Processing speed in Gflops

Here are some links with further information:

Main page: <http://www.netlib.org/benchmark/hpl/>

Description of the algorithm:

<http://www.netlib.org/benchmark/hpl/algorithm.html>

FAQ: <http://www.netlib.org/benchmark/hpl/faqs.html>

Tuning: <http://www.netlib.org/benchmark/hpl/tuning.html>

Appendix C - AMD Proposal**University Funding Request****Date:** September 12, 2001**University:** University of Central Florida (UCF)**Program Name:** OPCODE**Program Overview:**

Provide a brief description of the program for which you are requesting support.

The OPCODE Cluster project is really a collaboration between several different projects and funding sources. Chief among these is the original OPCODE project funded by the Air Force Office of Scientific Research (AFOSR) under the Defense/University Research Initiative Program. AFOSR has provided approximately \$217,000 in funds strictly for use in equipment purchases to build two large computing clusters at UCF. In addition, partial matching funds have been provided by the UCF Office of Sponsored Research (\$27,000), the UCF School of Electrical Engineering and Computer Science (SEECs) (\$13,500), and the Institute for Simulation and Training (IST) (\$13,500). The OPCODE Cluster project is an outgrowth of an earlier cluster computing project jointly funded by the Army Simulation and Training Instrumentation Command (STRICOM) and UCF, which resulted in the construction of our original 16 node, 32 processor cluster. STRICOM's Advanced Tactical Engagement Simulation Program and UCF are providing an additional \$27,000 for the new OPCODE project, above and beyond funds cited previously. Dr. Guy Schiavone is Principal Investigator these cluster projects, and is also technical lead on another UCF cluster project funded internally by UCF SEECs that has resulted in the recently completed "SCEROLA" cluster, employing 128 900 MHz Athlon Thunderbird processors.

System:

Describe the cluster system you intend to build.

The OPCODE project will result in the construction of two 192-processor computing clusters, one located at IST and the other at SEECs UCF. These clusters will employ a dual channel fast Ethernet network, linked together over a high-speed, fully stacked switch. The SEECs-OPCODE cluster will be linked on the same switch with the newly completed 128-processor SEECs cluster cited above. At each node, we plan to employ dual AMD processors on a Tyan

"Tiger" K7 motherboard, with 1 GB of DDR RAM, and 20 Gb local disk storage per node. The Linux operating system will be used, and a large number of supporting packages will be installed to support general needs of research and education in distributed and parallel computation. To support video and audio processing capabilities, an alternative IEEE 1394 network will also be available. We have contributed extensively to developing IEEE 1394 drivers for Linux, and one of our patches is included in the Linux 2.4.x kernel.

Application:

Why should AMD be particularly interested in your proposed system, and what it can offer to industry.

The twin OPCODE clusters being developed will be one of the first to employ dual AMD processors, and will thus demonstrate the power and cost efficiency of the new AMD multiprocessing capabilities. Another novel aspect of these clusters is the alternative IEEE 1394 network that will be used for video and audio processing. For example, use of the 1394 network is incorporated in a pending proposal to DARPA, on the topic of Voice Recognition in Background Noise, a surveillance technology particularly relevant in light of recent terrorist attacks on our country. Another unique proposed application of the cluster is in the area of Computer Generated Force Simulations for the Army OneSAF program. The use of cluster computing in this area will result in the generation of a large number of battlefield entities with more sophisticated behaviors than previously possible. We plan an aggressive benchmarking schedule, and fully expect these computers to rank on the list of Top 500 most powerful computers in the world. In addition, we are developing graduate and undergraduate level course curricula in the area of cluster computing.

Benefits

Describe what AMD gains by participation. Why is this a good investment?

The OPCODE project is a large, high-visibility project that will result in a number of unique applications of cluster computing. We believe that our experience and vision in the area of cluster computing will result in a number of notable accomplishments, and we plan to aggressively publicize these accomplishments, both in the usual academic circles, and in the private sector. Our clusters will serve as a showcase for the success of AMD multiprocessing. We are strong proponents of AMD technology, and would welcome the opportunity to share the credit with AMD for our upcoming successes in the area of cluster computing.

Permissions

Will AMD have the University's permission to post your story on our Web site, in printed collateral, AMD customer presentations, and in materials sent to the media.

Yes, most definitely.

Proposal for 2001 ('01-'02):

Provide details on funding recommendation.

Our current budget is stretched, and will not accommodate the use of the new AMD Athlon 4 (Palomino) processors that are recommended for multiprocessing by AMD. We are requesting a donation of 192 AMD Athlon 4 (Palomino) processors from AMD to be used in our two OPCODE clusters. We will purchase the remaining 192 processors using our existing project funds.

References:

Points-of-Contact:

Dr. Guy Schiavone
Assistant Professor of Computer Engineering
School of Electrical and Computer Engineering
University of Central Florida
(407)882-1300
guy@ist.ucf.edu